Graphics Internetworking: Bottlenecks and Breakthroughs

Don Brutzman
Code UW/Br, Naval Postgraduate School
Monterey California 93943 USA
831.656.2149 voice, 831.656.3679 fax
brutzman@nps.navy.mil

http://web.nps.navy.mil/~brutzman/vrml/breakthroughs.html

Chapter 4, *Digital Illusions*, Clark Dodsworth editor,
Addison-Wesley, Reading Massachusetts, 1997, pp. 61-97.

1. Introduction and Motivation. Although networking is considered to be "different" than computer graphics, network considerations are integral to large-scale interactive three-dimensional (3D) graphics. Graphics and networks are now two interlocking halves of a greater whole: distributed virtual environments. New capabilities, new applications and new ideas abound in this rich intersection of critical technologies. Our ultimate goal is to use networked interactive 3D graphics to take full advantage of all computation, content and people resources available on the Internet.

Network breakthroughs repeatedly remove bottlenecks and provide new opportunities. A pattern appears as we attempt to scale up in capability and capacity without limit: every old bottleneck broken reveals another. Understanding the bottlenecks, the corresponding solutions and potential upper bounds to growth permits us to develop effective networked graphics. When we overcome current bottlenecks, "effectively networked graphics" will simply mean "applications."

Internetworked graphics can be examined from the perspectives of connectivity, content, interaction, economics, applications and personal impacts. Internetworking refers to the ability to seamlessly interconnect multiple dissimilar networks globally. Connectivity has numerous dimensions including capacity, bandwidth, protocols and the many-to-many capabilities of multicasting. Content equals the World Wide Web and includes any type of information, dataset or data stream that might be used in the graphics environment. Interaction implies minimal latency, a sense of presence, and the ability to both access and modify content. The economics of networked graphics environments is developing rapidly and principal forces can be identified. Applications drive infrastructure development and are the most exciting part of networked graphics. Finally, the personal impacts which accompany these developments range from trivial to profound as high-quality interactive internetworked computer graphics become the norm on all computers.

2. Connectivity. There are many component pieces to the network connectivity puzzle. We will examine physical connectivity, network layers, the Internet Protocol suite, multicast distribution, the Multicast Backbone (MBone), enabling individuals, and vertical integration. An extensive overview of many networking topics is necessary to provide fundamental technical background. An understanding of each of these components is needed if we are to successfully implement highly shared graphics applications.

Physical connectivity to the Internet is a prerequisite to internetworked computer graphics. Other types of connections do not scale. There are several types of Internet access, roughly corresponding to direct providers, online services and bulletin board systems (BBSs). Foremost is "direct" connectivity where a user's computer is connected via a telephone modem, local-area network (LAN) or other link to an Internet service provider. The native communication protocol used across these direct links and the global Internet is the Internet Protocol (IP). Designed for minimum complexity and maximum capability, IP compatibility and direct Internet connectivity are essential for any cutting-edge application that uses "anybody anywhere anytime any-kind" connectivity. Lesser degrees of connectivity are possible via gateways when indirectly connected through commercial services such as America On-Line (AOL) or CompuServe. Even lower degrees of connectivity like file transfer and store-and-forward e-mail are possible through dial-up BBSs. To achieve the maximum breakthroughs in networked computer graphics, direct Internet connectivity is essential. We assume such connectivity throughout this chapter.

Network protocols and layered models. To integrate networks with large-scale graphical virtual environments, we invoke underlying network functions from within applications. Figure 1 shows how the seven layers of the well-known Open Systems Interconnection (OSI) standard network model generally correspond to the effective layers of the Internet Protocol (IP) standard. Functional characteristic definitions of the IP layers follow in Figure 2. These diagrams and definitions are merely an overview to help illustrate the logical relationship and relative expense of different network interactions. OSI and IP are often viewed as competing approaches to networking. A good overview of relative strengths and weaknesses of OSI versus IP is presented in [Malamud 92]. In practice OSI-related protocols are not of interest due to typically greater overhead and closed standards. For graphics internetworking, IP compatibility is essential for reasons of performance, open specifications, development flexibility and overwhelmingly strong global deployment trends.

In general, network operations consume proportionately more processor cycles at the higher layers. We must minimize this computational burden to reduce latency and maintain real-time responsiveness. Two protocols are currently available in the transport layer: one ensures reliable delivery, the other makes no guarantees (therefore requiring less overhead). Information transfer must use either reliable connection-oriented Transport Control Protocol (TCP) or nonguaranteed-delivery connectionless User Datagram Protocol (UDP). We choose between the two protocols depending on the criticality, timeliness and cost of imposing reliable delivery on the particular stream being distributed. Understanding the precise characteristics of TCP, UDP and related protocols helps the virtual-world designer understand the strengths and weaknesses of each network tool employed. Since internetworking considerations impact all components in a large virtual environment, understanding network protocols and application hooks is essential for virtual-world designers [Internic 95] [Stallings 94] [Comer 91] [Stevens 90].

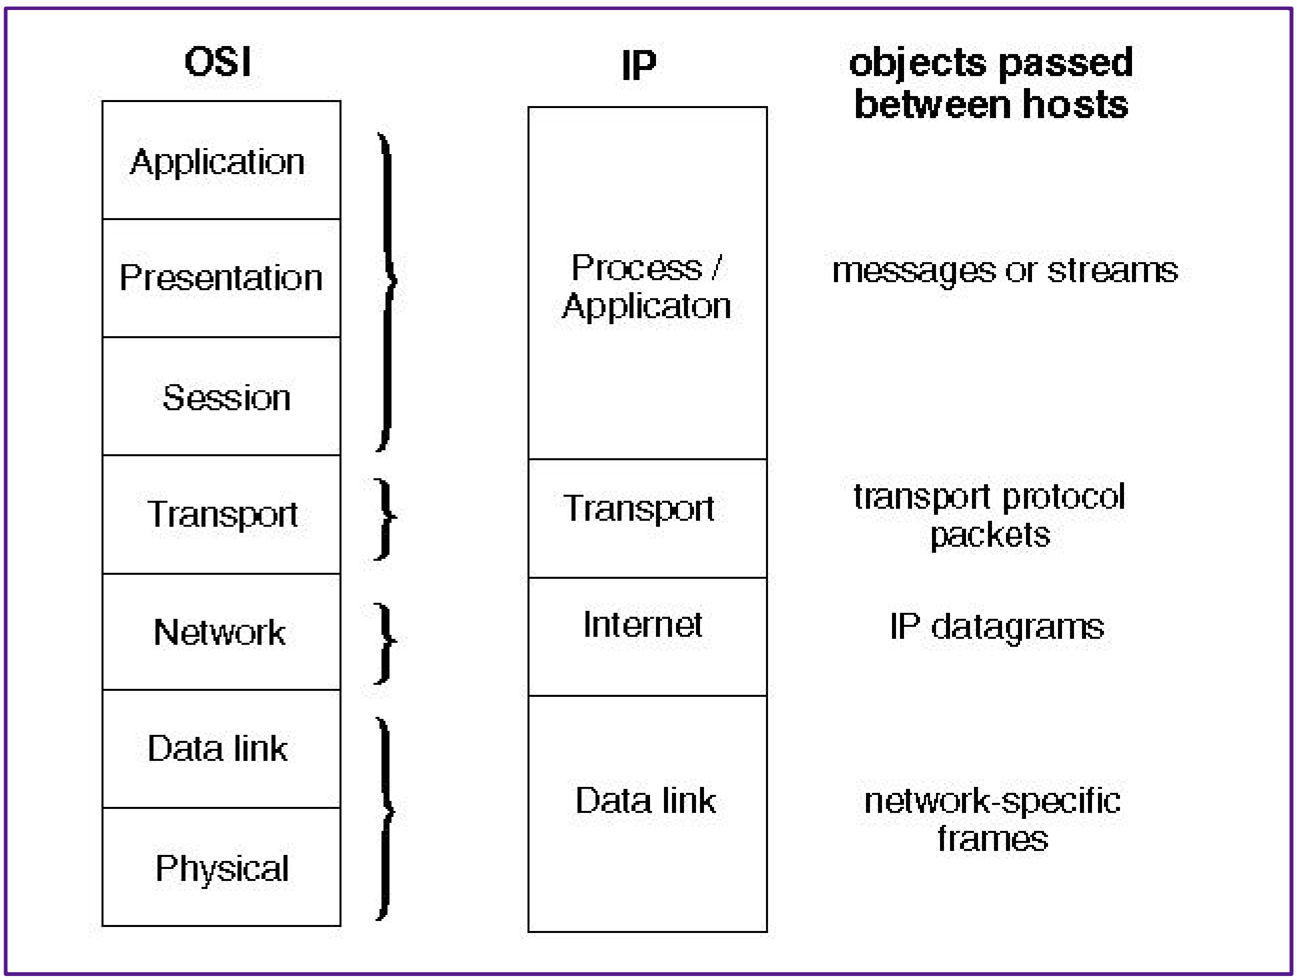

Figure 1. Correspondence between OSI and IP protocol layer models, and objects passed between corresponding host layers.

Internet Protocol (IP). Although the protocols associated with internetworking are diverse, there are some unifying concepts. Foremost is "IP on everything," the principle that every protocol coexists compatibly within the Internet

Protocol suite. The global reach and collective momentum of IP-related protocols makes their use essential. Current IPv4 and next-generation IPv6 [Bradner 95] protocols can be run over a complete variety of electrical, radio-frequency and optical physical media.

Examination of protocol layers helps clarify network issues (Figure 2). The lowest layers (Physical, Data Link) are reasonably stable, with a huge installed base of Ethernet and FDDI systems, augmented by the rapid development of wireless and broadband ISDN solutions such as Asynchronous Transfer Mode (ATM). Compatibility with the Internet Protocol (IP) suite is assumed. The middle transport-related layers (Internet, Transport) are a busy research and development area. Real-time reliability, quality of service, and other capabilities are evolving rapidly due to competing market forces, both academic and commercial. Lower- and middle-layer bottlenecks are mostly solvable using available techniques. The major bottlenecks to large-scale graphics internetworking are at the application layer.

Process/Application Layer. Invokes TCP/IP services, sending and receiving messages or streams with other hosts. Delivery can be intermittent or continuous.

Transport Layer. Provides host-host packetized communication between applications, using either reliable-delivery connection-oriented TCP or unreliable-delivery connectionless UDP. Exchanges packets end-end with other hosts.

Internet/Network Layer. Encapsulates packets with an IP datagram that contains routing information, receives or ignores incoming datagrams as appropriate from other hosts. Checks datagram validity, handles network error and control messages.

Data Link/Physical Layer. Includes physical media signaling and lowest level hardware functions, exchanges network-specific data frames with other devices. Includes capability to screen multicast packets by port number at the hardware level.

Figure 2. Summary of Internet Protocol (IP) suite layer functionality.

Multicast. Multicast packet distribution is best understood in comparison with two other packet distribution mechanisms: unicast and broadcast. Unicast is point-to-point communication, commonly occurring whenever an e-mail message is sent or a Web browser connects to a home page. Only the recipient host and intermediate routers need to spend computational cycles on a unicast packet. One-to-many unicast communications are accomplished by multiple streams corresponding to every recipient. While this is not a problem when sending an e-mail message to multiple recipients, it is clearly an expensive proposition when sending a high-bandwidth stream such as graphics or video.

Broadcasting is at the opposite end of the spectrum from unicast. Broadcast messages reach every host on a local-area network and demand a response from each application or operating system. This is extremely inefficient at high bandwidths. Furthermore, each broadcast packet must be read and processed by every host on a network. Broadcast is typically prohibited from passing across routers to prevent such packets from propagating everywhere. Figure 3 shows typical unicast and broadcast packet stream paths.

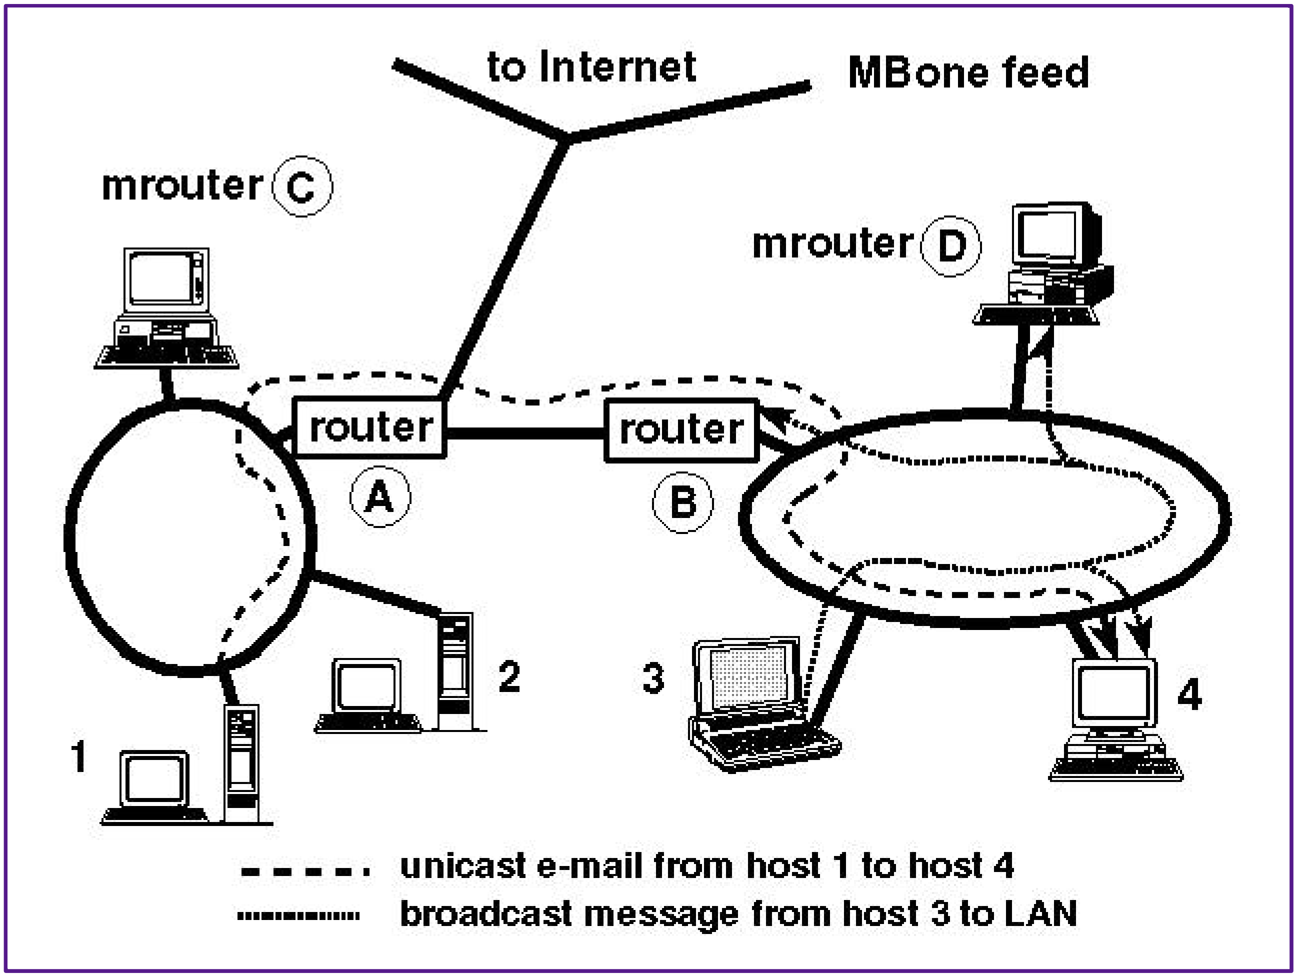

Figure 3. Typical unicast and broadcast packet stream paths.

Multicast packet functionality lies between unicast and broadcast. A single multicast packet can touch every host on a LAN without duplication. However each host can subscribe or ignore multicast at the hardware level by informing the network interface card which multicast channels (if any) to monitor. This feature is extremely important since a single high-bandwidth stream can still reach an arbitrary number of hosts, but computational load is only seen on hosts which explicitly subscribe to the multicast channel. Multicast packets are usually prevented from traversing routers to prevent overloading the Internet with arbitrary propagation of such streams. Figure 4 illustrates this behavior. To summarize, the key network characteristics which distinguish multicast from unicast and broadcast include:

- individual multicast packets can be read by every host workstation on a LAN
- elimination of duplicate streams minimizes bandwidth requirements
- workstations can screen unwanted multicast packets at the hardware level, eliminating unnecessary computational burden on applications

---

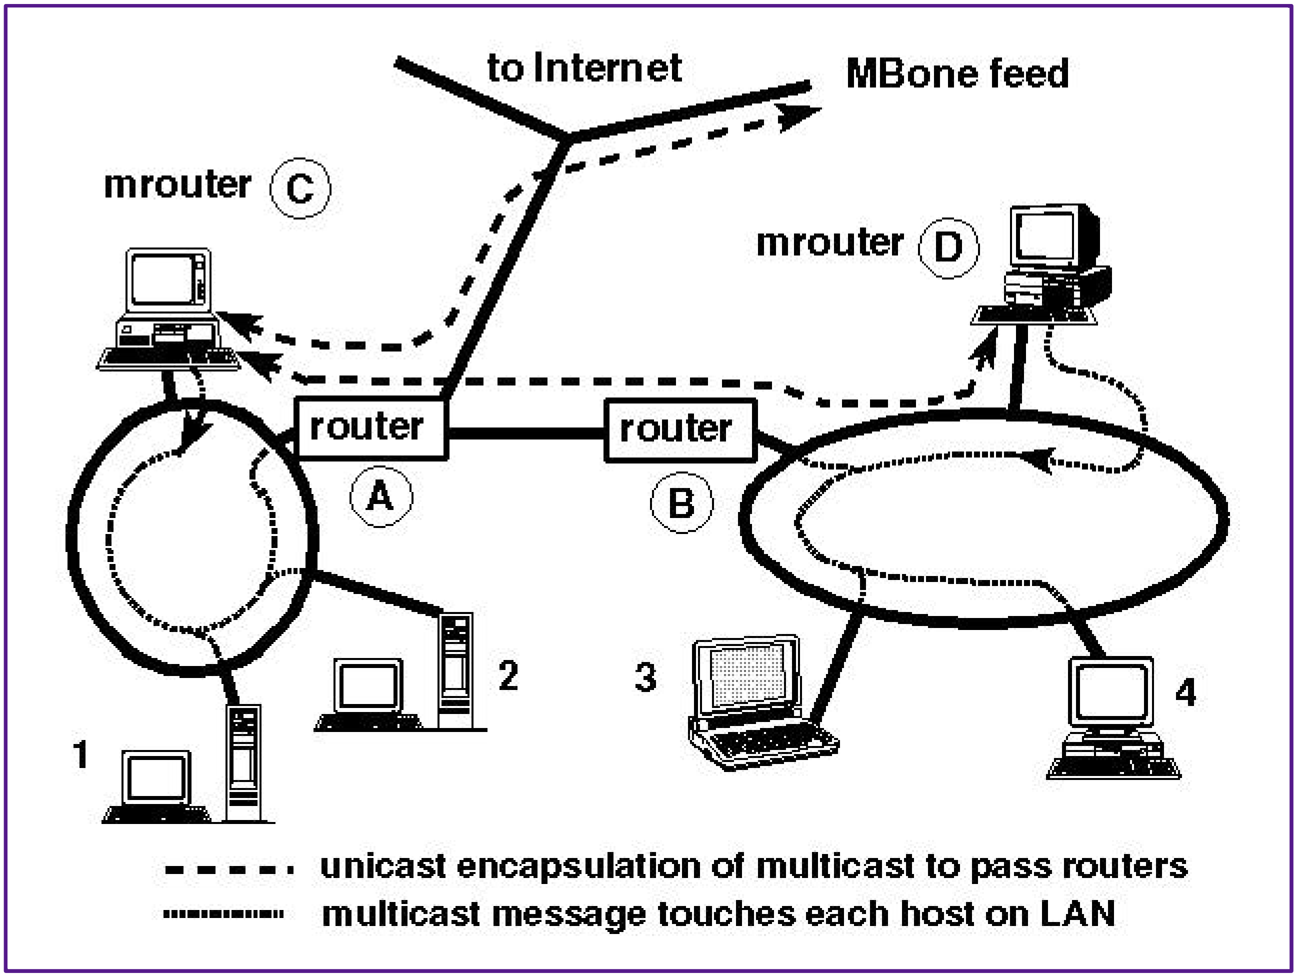

Figure 4. Example of multicast routers connecting multicast streams encapsulated by unicast (in order to bypass router prohibitions on multicast). After stripping unicast headers, multicast streams are then available to each host on the LAN. Direct multicast support in routers has recently become available. Such multicast routes between LANs comprise the Multicast Backbone (MBone).

Multicasting has existed for several years on local-area networks such as Ethernet and Fiber Distributed Data Interface (FDDI). Now, with Internet Protocol multicast addressing at the network layer, group communication can be established across the Internet. IP multicast addressing is an Internet standard developed by Steve Deering [Deering 89]. Categorized officially as an IP Class D address, an IP multicast address is merely a set of reserved IP addresses in the range (224.0.0.0 thru 239.255.255.255). Multicast is supported by numerous workstation vendors including SGI, Sun, DEC and HP.

Unicast audio/video applications connecting groups of people across the Internet can be harmful because they waste bandwidth and do not scale. Recent examples include PC and Macintosh implementations of CU-SeeMe [Cogger 95]. Multicast applications have long been expected for Intel and Apple processors, but first multitasking must be better implemented in personal computer operating systems. Current audio/video applications for personal computers are typically unicast and need to be modified to support multicast before widespread use of arbitrarily shared data streams can be expected for all computers on the Internet.

Multicast Backbone (MBone). The MBone is one of the Internet's most interesting capabilities and is used for distribution of live audio, video and other packets on a global scale. MBone is a virtual network because it shares the same physical media as the Internet. It uses a network of routers (mrouters) that can support multicast. These mrouters are either upgraded commercial routers, or dedicated workstations with modified operating system kernels running in parallel with standard routers. It was named by Steve Casner and originated from an effort in 1992 to multicast audio and video from meetings of the Internet Engineering Task Force (IETF). The fact that the MBone works at all is pretty remarkable since the original multicast backbone code was a student experiment. Encapsulated multicast streams sidestep regular routers through handcrafted tunnels across the Internet. Thus the

MBone exemplifies the best of the hacker ethic on a large scale. This pioneer spirit is still visible today as we cooperatively transition to a more carefully engineered backbone service that includes widespread multicast service.

Two things make multicasting feasible on a world-wide scale: installation of high-bandwidth Internet backbone connections, and widespread availability of workstations with adequate processing power and built-in audio capability. Today hundreds of researchers use MBone to develop protocols and applications for group communication. Multicast provides one-to-many, several-to-many and many-to-many network delivery services. A one-to-many example might be live transmission of an academic conference. A several-to-many event might be a panel interview between different locations attended by many dispersed people. A many-to-many example might be a distributed battlefield exercise with 500 active players. Thus multicast is useful for a variety of videoconferencing, audio and multiplayer events where numerous hosts need to communicate simultaneously. Network concepts underlying MBone, bandwidth considerations, application tools, MBone events, interesting MBone uses, and how to connect to the MBone are described in "MBone Provides Audio and Video across the Internet" [Macedonia, Brutzman 94].

Lost MBone packets stay lost. Since multicast currently uses only the User Datagram Protocol (UDP) and not the Transport Control Protocol (TCP) of the IP suite, multicast streams are connectionless "unreliable" service. No setup is required, no acknowledgements are used, and no guarantee of delivery exists for this type of "best-effort" service. This is ordinarily a good thing for most real-time information streams (such as audio and video) since it avoids delivery bottlenecks and unwanted overhead. Furthermore it doesn't matter what kind of data is sent via multicast IP packets. Several researchers are experimenting with "reliable multicast" transport protocols which try to achieve a different balance by gaining occasional retransmission benefits without unacceptable acknowledgement overheads [Schulzrinne 96]. These considerations are critical when scaling to arbitrarily large numbers of simultaneously interacting users on the network. A breakthrough in this bottleneck will enable new global communication paradigms for widely distributed shared applications.

The keys to understanding MBone constraints are capacity and bandwidth. The reason a multicast stream is bandwidth-efficient is that one packet can touch all workstations on a network. A 64-kilobit per second (Kbps) video stream typically provides 3 to 5 frames per second of quarter-size NTSC video. Each multicast stream uses the same bandwidth whether it is received by 1 workstation or 20. That is good. However, that same multicast packet is ordinarily prevented from crossing network boundaries such as routers. The reasons for this current restriction are religious and obvious from a networking standpoint. If a multicast stream that can touch every workstation jumps from network to network without controls, the entire Internet might quickly become saturated by such streams. That scenario is disastrous, so group controls are necessary. The MBone controls multicast packet distribution across the Internet in several ways. The topological "lifetime" of multicast packets is constrained to limit the number of allowed mrouter hops, and automatic pruning/grafting algorithms adaptively restrict multicast transmission to active participants. Such controls can prevent most (but not all) global MBone problems.

Responsible daily use of the MBone network merely consists of making sure you don't overload your local or regional bandwidth capacity. Appropriate bandwidth values initially seem obscure but daily practice quickly makes such figures intuitive. MBone protocol developers are successfully experimenting with automatically pruning and grafting subtrees, but for the most part MBone uses thresholds to truncate broadcasts to the leaf routers. The truncation is done by the setting the "time-to-live" (ttl) field in a packet that is decremented by at least one each time the packet passes though an mrouter. A ttl value of one prevents a multicast packet from escaping a LAN. An initial ttl value of 16 typically limits multicast to a campus, as opposed to values of 127 or 191, which might send a multicast stream to every subnet on the MBone (currently over 20 countries). Sometimes we decrement ttl fields by intentionally large predetermined values to limit multicasts to sites and regions.

These issues may sound challenging but are relatively uncomplicated in practice. Some personal technical proficiency is important because use of the MBone can have a major impact on shared network performance. For example, default video and audio streams consume about 200 Kbps of bandwidth, over 10 percent of a 1.5-Mbps T1 line (a common site-to-site link on the Internet). Several simultaneous high-bandwidth sessions might easily saturate network links and routers. The number of active users on current MBone channels is usually small enough that consensus and planning are able to effectively share this limited global resource. These low numbers will change soon. Planning, consensus and careful utilization of physical and logical network resources will remain essential components of live large-scale global internetworking.

> "It is not every day that someone says to you, 'Here is a multimedia television station that you can use to broadcast from your desktop to the world.' These are powerful concepts and powerful tools that extend our ability to communicate and collaborate tremendously. They have already changed the way

people work and interact on the net." [Macedonia, Brutzman 94]

Enabling individuals. More and more often, the cost of admission for participation in global internetworking is a willingness to learn and participate. The cost of equipment is relatively low and the price of bandwidth continues to become more affordable in most countries. As physical infrastructure becomes more prevalent, physical connectivity becomes easily available for many people. Public domain software for almost any purpose can be obtained freely over the Internet. Frequently asked question lists (FAQs), mailing lists and even online books provide amazing depth on almost every subject related to networking and graphics. Individuals are now able to do many things that previously required large groups and large budgets to accomplish. This trend is a tremendous equalizer which lets students, researchers and companies become effective players in a global arena.

Vertical integration. For many years there have been walls separating communities of users corresponding to the computer architecture or operating system used. Supercomputers, mainframes, workstations and personal computers were so dissimilar that commonality between them was practically nonexistent. Networked applications have changed this situation completely. Once-isolated islands are now complementary, interdependent and practically seamless.

Initial rudimentary examples of interconnectivity are e-mail, file transfer protocol (ftp) and remote terminal protocol (telnet). The Internet Protocol suite enables standardized hostnames, host numbers and Uniform Resource Locators (URLs), which in turn convert loosely connected hardware into a shared global address space. Sharable applications such as Mosaic, Netscape, distributed interactive simulation and the MBone tools show that equivalent functionality can be expected on any computer platform, from the supercomputers to PCs. The Hypertext Markup Language (HTML) and the Web now provide a universal interface to any computational process, a fact surprisingly missed by most interface designers. This new interoperability between all computers means that applications designers can think in terms of vertical integration. Thousands of simultaneous users with networked personal computers can now query, drive and evaluate computationally intensive graphics scenes rendered by powerful computers at arbitrarily distant locations. Such scenarios have been described theoretically for years. While theoretical discussions continue, the practical capability has arrived. Together, all the countless machines interconnected by the Internet are our new supercomputer.

Here are the success metrics for vertical interoperability: "Will the application run on my supercomputer?" *Yes*. "Will it run on my Unix workstation?" *Yes*. "Will it also run on my Macintosh or PC?" *Yes*. This approach has been shown to be a practical (and even preferable) software requirement. Vertical interoperability approaches are best supported by open nonproprietary specifications developed by responsive standardization groups such as the Internet Engineering Task Force (IETF).

Broad interoperability and Internet compatibility are essential. Closed solutions are dead ends. In order to achieve broad vertical integration, proprietary and vendor-specific hardware need to be avoided. Videoteleconferencing systems are an example of a market fragmented by competing and incompatible proprietary specifications. In the area of new connection-oriented cell-based services such as Asynchronous Transfer Mode (ATM) and Integrated Services Digital Network (ISDN), some disturbing trends are commonplace. Supposedly standardized protocol implementations often do not work as advertised, particularly when run between hardware from different vendors. Effective throughput is often far less than maximum bit rate. Latency performance is highly touted and rarely tested. Working applications that use these services are difficult to find. Corresponding network operating costs are often hidden or ignored. Perhaps worst of all, long-haul services such as ATM may not be fully interoperable with the Internet Protocol and some core functionality (such as many-to-many multicast) may not be feasible. Application developers are advised to plan and budget for lengthy delays and costly troubleshooting when working with new services.

Connectivity conclusions. We have seen that connectivity has many facets. A technical overview of these concepts was necessary in order to understand what is feasible. Our technical foundation is now strong enough to consider broader concepts of content, interaction, economics, applications and personal impacts.

3. Content and the World Wide Web. Connectivity isn't worth much unless there's information of value (content) to exchange. While human and artificial users have created a wide variety of content for many years, only recently has it been available on a global scale. Many proposals have been made regarding the optimum way to achieve a global database, but only one methodology is a contender: the Web.

If the Internet is our global supercomputer, then the Web is our global database. The Web has been defined as a "wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe of documents" [Hughes 94]. Fundamentally, the Web combines a name space consisting of all visible information stores on the Internet with a broad set of retrieval clients and servers, all of which can be connected by easily defined Hypertext Markup Language (HTML) hypermedia links. This globally accessible combination of media, client programs, servers and hyperlinks can be conveniently utilized by humans or autonomous entities. The Web continues to fundamentally shift the nature of information storage, access and retrieval, best explained by its originator [Berners-Lee 94]. Information content is the crucial component underlying most 3D graphics renderings. Information context is the key quality provided by the universal naming scheme and simplicity of HTML. Anyone can now add any content and any context they might want.

Current Web capabilities are easy to use despite rapid growth and change. An extension mechanism allows new content types and corresponding viewers to be easily integrated into existing browsers. The relative simplicity of HTML has been essential to the rapid growth of the Web. People add their own content because it's easy to do. Convention and common use have standardized many content formats, and ease of use has been the principal engine of success.

The ubiquitous availability of content via the Web has massive implications that we are only beginning to perceive. It's one thing to be aware that a Library of Congress exists somewhere, with most information sources that might ever be wanted. It's another to have such an information archive close at hand for easy access, regardless of your physical location [Library 95]. It's a further step to think that such an archive might be rapidly searchable and conveniently retrievable on demand. Until recently it was unprecedented to think that anyone anywhere might contribute to such a store without formal constraint. Now such capabilities are common. Future content challenges include easy connections for datasets, databases and analytic models to the Web.

These ideas are a series of comprehensible steps, but together they imply amazing capabilities. If knowledge is power, then these capabilities are powerful indeed. It is now possible for any question on nearly any subject to quickly provoke more information than is humanly possible to absorb. Eventually "everybody's everything" can be online. It seems there is now an ocean of information surrounding us. We can swim around the world or drown trying to drink it all in. How we deal with it is our choice.

While the full implications of universal access to massive content are still emerging, it's clear the genie can't be stuffed back into the bottle. We must deal with all this content. Such a challenge is a good match for interactive 3D computer graphics. Since graphics techniques can portray large amounts of information in a meaningful way, realistic rendering and scientific visualization techniques are logical choices for dealing with massive content. Effective 3D user interfaces are also needed for navigating the richness of these worlds. Interfaces must feel intuitive to both novice and expert, particularly if access and interaction are to scale up to match amazingly large volumes of content.

There are two more important points related to networking and the Web: improved interaction and importance of an open standards process. The Hypertext Transfer Protocol (http) which is used for most Web-based interactions typically provides a client-server relationship; a user can push on a Web resource and get a response, but there are no openly standardized mechanisms for a Web application to independently push back at the user. Netscape has independently defined and implemented server-push and client-pull techniques that permit repeated queries and replies [Netscape 95]. An example of "push" is when a server can send repeated updates down an open connection to refresh a client display. An example of "pull" is automatic repeated queries by a client at intervals matching data refresh cycles on the server. These methods for improving interaction on the Web are steps in the right direction. Better interaction mechanisms will enable new types of Web applications.

Netscape preempted the HTML standards process by unilaterally proposing and incorporating push/pull capabilities. Although server push and client pull are useful concepts likely to be included in forthcoming versions of the HTML specification, unilateral implementations can fragment a common standard into incompatible, inadequate variants. Open standards review means that any specification can be widely scrutinized in detail. Meanwhile (in the IETF at least) two or more independent implementations are needed to validate correct performance prior to acceptance as a standard. Patience and cooperation are essential since coopting standardization through market share creates far more bottlenecks than breakthroughs. Considering this particular technical example, as the demands of interactive client/server queries and replies grow in scale, many of the bandwidth and scaling issues explored in the MBone will reappear. Since cooperative standards development has been essential for the success of the MBone, similar cooperative standards-based efforts will continue to be essential as we scale up the Web.

4. Virtual Environments (VEs). The scope of virtual environment development is so broad that it can be seen as an inclusive superset of all other global information infrastructure applications. Virtual environments and virtual reality applications are characterized by human operators interacting with dynamic world models of truly realistic sophistication and complexity [Zyda 93] [Durlach 95]. Current research in large-scale virtual environments can link hundreds of people and artificial agents with interactive 3D graphics, massive terrain databases, global hypermedia and scientific datasets. Related work on teleoperation of robots and devices in remote or hazardous locations further extends the capabilities of human-machine interaction in synthetic computer-generated environments. VE construction can include concepts and components from nearly any subject area. The variety of desired connections between people, artificial entities and information can be summarized by the slogan "connecting everyone to everything." As diversity and detail of virtual environments increase without bound, network requirements become the primary bottleneck.

The most noticeable characteristic of virtual environments is interactive 3D graphics, which ordinarily involves coordinating a handful of input devices while placing realistic renderings at fast frame rates on a single screen. Networking can connect virtual worlds with realistic distributed models and diverse inputs/outputs on a global scale. Graphics and virtual-world designers interested in large-scale interactions can now consider the world-wide Internet as a direct extension of their computer. A variety of networking techniques can be combined with traditional interactive 3D graphics to provide almost unlimited connectivity. Experience shows that the following services are essential for virtual world communications: reliable point-to-point communications, behavior interaction protocols such as the IEEE standard Distributed Interactive Simulation (DIS) protocol, World Wide Web (WWW) links, and multicast channels for bandwidth-efficient many-to-many communications. Further examination of behavior-based interaction protocols is next presented using DIS as an example methodology.

Distributed Interactive Simulation (DIS). The DIS protocol is an IEEE standard for logical communication among entities in distributed simulations [IEEE 93] [DIS 94]. Although initial development was driven by the needs of military users, the protocol formally specifies the communication of physical interactions by any type of physical entity and is adaptable for general use. Information is exchanged via protocol data units (PDUs) which are defined for a large number of interaction types.

The principal PDU type is the Entity State PDU. This PDU encapsulates the position and posture of a given entity at a given time, along with linear and angular velocities and accelerations. Special components of an entity (such as the orientation of movable parts) can also be included in the PDU as articulated parameters. A full set of identifying characteristics uniquely specifies the originating entity. Several complementary dead-reckoning algorithms permit computationally efficient projection of entity posture by listening hosts. Dozens of additional PDU types are defined for simulation management, sensor or weapon interaction, signals, radio communications, collision reporting and logistics support. Most of these highly specialized "garbage PDUs" are of questionable utility for most VE applications. What is essential is correct (or at least "close-enough") physics, real-time realism and adequate global consistency.

Here is an example showing how the DIS entity state PDU properly enables networked distribution of realistic physics-based behavior. In 1992 a jet pilot at the Naval Postgraduate School developed a real-time model of high-performance aircraft response using parameterized flight coefficients and quaternion mathematics [Cooke 92]. It turns out that tactical aircraft are inherently unstable. Jet pilots require years of training and onboard computers are continuously making compensatory adjustments to flight surfaces. Qualified fighter pilots are pleased with this quaternion-based model since the simulated jet reacts realistically and "flies" much like a true jet. Unfortunately, naive operators tend to crash when flying this model (and presumably might experience similar difficulties in a real jet). Substituting flight coefficients that represent a more benign aircraft like a Piper Cub made operation much easier for casual users. The significant fact related to this substitution is that the underlying DIS model requires no changes. Regardless of whether packet update rates necessary for smooth motion average once per 5 seconds or 20 times per second, the DIS Entity State PDU dead reckoning algorithms are able to properly track position and orientation. This maximizes precision, maintains proper global consistency, and minimizes network loading.

Of further interest to virtual-world designers is an open-format Message PDU. Message PDUs enable user-defined extensions to the DIS standard. Such flexibility coupled with the efficiency of Internet-wide multicast delivery permits extension of the object-oriented message-passing paradigm to a distributed system of essentially unlimited scale. Free-format DIS Message PDUs might provide general message-passing connectivity to any information site on the Internet, extended by use of network pointer mechanisms that already exist for the World Wide Web. This is a promising area for future work.

VE grand challenges. The most important "grand challenges" of computing today are not large, static gridded simulations such as computational fluid dynamics or finite element modeling. Similarly, traditional supercomputers are not necessarily the most powerful or significant platforms because adding hardware and dollars to incrementally improve existing expensive computer designs is a well-understood exercise. What's more challenging and potentially more rewarding is the interconnection of all computers in ways that support global interaction of people and processes. In this respect, the Internet is the ultimate supercomputer, the Web is the ultimate database, and any networked hardware in the world is a potential input/output device. Large-scale virtual environments attempt to simultaneously connect many of these computing resources in order to recreate the functionality of the real world in meaningful ways. Network software is the key to solving VE grand challenges.

Large-scale virtual-world internetworking. Four key communication methods are necessary for large-scale virtual-world internetworking: light-weight messages, network pointers, heavy-weight objects and real-time streams (Figure 5). For each of these four methods, bandwidth and latency must be carefully considered. The following networking elements can be combined to implement these four methods.

- Distribution of virtual-world components using point-to-point sockets can be used for tight coupling and real-time response of physics-based models.
- DIS is a well-tested behavior protocol that enables efficient live interaction between multiple entities in multiple virtual worlds.
- A wide variety of alternative behavior protocols are being developed and evaluated for large-scale virtual-world interaction.
- Hypermedia servers and embedded Web browsers provide virtual worlds global access to pertinent archived images, video, graphics, papers, datasets, software, sound clips, text or any other computer-storable media, both as inputs and outputs.
- Multicast protocols permit moderately large real-time bandwidths to be efficiently shared by an unconstrained number of hosts.
- MBone connectivity permits live distribution of graphics, video, audio, DIS and other streams world-wide in real time.

Together these example components provide the core functionality of light-weight messages, network pointers, heavy-weight objects and real-time streams. Integrating these network tools in virtual worlds produces realistic, interactive and interconnected 3D graphics that can be simultaneously available anywhere [Brutzman 94] [Macedonia 95a, b].

Light-weight Interactions. Messages composed of state, event and control information as used in DIS Entity State PDUs. Implemented using multicast. Complete message semantics is included in a single packet encapsulation without fragmentation. Light-weight interactions are received completely or not at all.

Network Pointers. Light-weight network resource references, multicast to receiving groups. Can be cached so that repeated queries are answered by group members instead of servers. Pointers do not contain a complete object as light-weight interactions do, instead containing only a reference to an object.

Heavy-weight Objects. Large data objects requiring reliable connection-oriented transmission. Typically provided as an ftp or http response triggered by a network pointer request.

Real-time Streams. Live video, audio, DIS, 3D graphics images or other continuous stream traffic that requires real-time delivery, sequencing and synchronization. Implemented using multicast channels.

Figure 5. Four key communications components used in virtual environments.

Application layer interactivity. It is application layer networking that needs the greatest attention in preparing for the information infrastructure of the near future. DIS combined with multicast transport provides solutions for many application-to-application communications requirements. Nevertheless DIS is insufficiently broad and not adaptable enough to meet general virtual environment requirements. To date, most of the money spent on networked virtual environments has been by, for and about the military. Most remaining work has been in (poorly) networked games. Neither example is viable across the full spectrum of applications. There is a real danger that specialized high-end military applications and chaotic low-end game hacks will dominate entity interaction models. Such a situation might well prevent networked virtual environments from enjoying the sustainable and compatible exponential growth needed to keep pace with other cornerstones of the information infrastructure.

Next-generation behavior protocols. Successors to DIS are needed that are simpler, open, extensible and dynamically modifiable. DIS has proven capabilities in dealing with position and posture dead-reckoning updates, physically based modeling, hostile entity interactions and variable latency over wide-area networks. DIS also has several difficulties: awkward extendability, requiring nontrivial computations to decipher bit patterns, and being a "big" (i.e. complicated) standard. DIS protocol development continues via a large and active standards community. However, the urgent military requirements driving the DIS standard remain narrower than general virtual environment networking requirements. Smaller and more efficient behavior protocols are being investigated by several researchers. The current NPS approach for next-generation behavior communications is to develop a "dial-a-protocol" capability, permitting dynamic modifications to the protocol specification to be transmitted to all hosts during a simulation [Stone 96]. A dynamically adjustable protocol is necessary for interactively testing global and local efficiency of distributed entity interactions. Our research is motivated by a fundamental networking principle: realistic testing and evaluation are essential, because initial performance of distributed applications never matches expectations or theory. Thus we're intentionally focusing on mechanisms for finding bottlenecks and demonstrating breakthroughs.

Other interaction models. Many other techniques for entity interaction are being investigated, although not always in relation to virtual environments. Several breakthroughs are possible. Just like in the real world, many ways of interacting will be valuable in virtual environments.

- Common Gateway Interface (cgi) scripts are executed on the queried host, often using Hypertext Transfer Protocol (http) [Berners-Lee 94] query extensions as inputs.
- Java has sparked wide interest as a network-based behavior language. Java enables simple and secure passing of precompiled program object files (applets) for compatible multiplatform execution [Sun 95].
- Intelligent agent interactions are an active area of research driven by the artificial intelligence and user interface communities. Rule-based agents typically communicate via a message-passing paradigm that is a straightforward extension of object-oriented programming methods.
- Instead of exchanging messages, mobile programs called agents can transport themselves. Telescript is a promising language designed for safe and secure mobile agent interactions [White 95].
- Ongoing research by the Linda project uses a generalized message-passing construct called "tuples" for logical entity interaction, with emphasis on scaling up to indefinitely large sizes and participants [Gelernter 92].
- MUDs (multi-user dungeons) and MOOs (MUDs Object-Oriented) provide a powerful server architecture and text-based interaction paradigm well suited to support a variety of virtual environment scenarios [Curtis 94].
- Passing interpretable scripts over the network in conjunction with predeployed, precompiled source code has been widely demonstrated for the multiplatform 2D graphical user interface (GUI) Tool Control Language (Tcl) [Ousterhout 94].

Virtual Reality Modeling Language (VRML). The Web is being extended to three spatial dimensions thanks to VRML, a specification originally based on an extended subset of the SGI OpenInventor scene description language [Wernicke 94]. Key contributions of the initial VRML 1.0 standard were a core set of object-oriented graphics constructs augmented by hypermedia links, all suitable for scene generation by browsers on Intel and Apple personal computers as well as Unix workstations. The interaction model for 3D VRML browsers remains client-server, similar to most other Web browsers. 3D browsers are usually embedded in 2D browsers or are launched as helper applications when connecting to a 3D site. VRML specification development has been effectively coordinated by mail list, enabling consensus by a large, active and open membership [Pesce 94, 95] [Bell 94].

Difficult issues regarding real-time animation in VRML 2.0 (Carey 96) include entity behaviors, user-entity interaction and entity coordination. Currently animation scripts and experimental servers are used for such functionality. In order to scale to many simultaneous users, peer-to-peer interactions will be necessary in addition to client-server query-response. Although "behaviors" are not formally specified, VRML 2.0 provides local and remote scripting language hooks (i.e. an applications programming interface or API) to graphical scene descriptions. Dynamic scene changes can thus be stimulated by scripted actions, message passing, user commands or behavior protocols, implemented using either Java or complete VRML scene replacement. A great deal of experimentation is in progress. It appears clear that successful global solutions for VRML behaviors are possible that provide simplicity, security, scalability, generality and open extensions. The most successful experimental approaches to shared behaviors will likely establish employment conventions and may even inspire behavior specifications. Finally, as the demanding bandwidth and latency requirements of virtual environments begin to be exercised by VRML, some client-server design assumptions of the Hypertext Transfer Protocol (http) may no longer be valid. Users won't be satisfied with network mechanisms that fail to accommodate high-bandwidth information streams or break down after a few hundred players. A Virtual Reality Transfer Protocol (vrtp) will be needed to take advantage of available transport-layer functionality and overcome bottlenecks in http. Experimentation and quantitative evaluation will be essential to develop the next generation of code for diverse interentity virtual environment communications. To achieve these ambitious goals, it is necessary to create a Cyberspace Backbone (CBone) providing dedicated network resources for a virtual environment testbed.

Note that computer graphics applications now become more flexible than television. From a network perspective, prerendered computer graphics streams are equivalent to video streams. Behavior-based VRML will provide dynamic graphics scene delivery at relatively low bandwidths with viewpoints independently controlled by subscribing users. Effective deployment of globally networked graphics can complement television in fascinating ways.

5. Economics. Predicting economic developments in new technologies is a risky business. The following points are only a few of many changes, emphasizing major new developments that are likely to produce market impacts. Users and investors can draw their own conclusions.

- Marketplace providers may be far different from traditional companies. Individuals and small groups of people can be economic players, for example. GigaCorp Intergalactic UltraServices Ltd. might be one person working at home producing information products that are as good as (or better than) the products a large multinational corporation produces. Unrestricted access to information, open interest-group efforts and affordable high-quality software are key. These assets produce the new trump cards in creative product development: imagination, cooperation and rapid turnaround.
- Value-added becomes a market basis. At first it might seem that when so much is free, no one needs to buy anything. In practice, people are willing to pay when there's significant value-added relative to what's free. For example, numerous software products are available today in identical free and supported versions. Nevertheless today's innovative value-added feature is tomorrow's commonplace free feature. Thus ongoing innovation, responsiveness and improvement are necessary for software program survival.
- What is "free" in the computer software community is a sliding standard that tracks along with the exponential growth of the computer price/performance technology curve. Historically, companies charging for quality software have been first to market, with quality freeware versions following later. Programmers unwilling or unable to pay fees often create public domain software comparable to commercial versions. Current free availability of quality software development tools and code libraries has accelerated this growth of free software projects. Freeware programmers are often willing to port software to multiple architectures. Mailing lists, news groups, online hyperlinked documentation and software repositories encourage the growth of diverse and technically proficient user communities. High-quality online experience is usually available to open interest-group members (however informally) at all hours of the day or night. Open cross-platform public-domain (free) solutions frequently inspire profitable commercial software products. Both climb the cutting-edge-technology curve in tandem.
- Information dissemination is easier than ever. Content for mass access is no longer solely controlled by television networks, radio broadcasters and print publishers. As more people put information online and collaborate with groupware, traditional information sources become less necessary. Just as television didn't put radio out of business, expect to see evolutionary changes in information sources rather than elimination of other media. Open access to disparate media over the network fosters both competition and cooperation. These are democratizing and equalizing forces. Everybody will want access to everything, and networking makes it possible. Interactive 3D graphics with independent viewpoint, navigation and content already provides serious competition to television. New 3D "channels" are creating new businesses, arts and entertainments.
- It becomes increasingly difficult to preserve original sources as internetworked information grows exponentially. Digital libraries will become essential to provide continuity of access. Context may be supported through ever-more-sophisticated links, browsers and virtual environments, but secure storage of basic content will remain a fundamental requirement. A forward-looking overview of digital library efforts is given by [Garrett 95].
- Informal market analysis of an industry that undergoes a complete reorganization every year or two indicates that open methods based on responsive standards survive, while closed proprietary methods struggle and die, or get absorbed/killed/subsumed by dominant companies in a market (like Microsoft). Semiopen methods (like ISO OSI) that are overconstrained by standardization or restrictive rules fall behind the pace of the marketplace and become obsolete - in effect they appear proprietary.
- Security considerations affect all parts of the Internet, including multiuser virtual environments. Encryption, authentication and nonrepudiation are all feasible using a variety of currently available open and closed cryptographic standards. Public key cryptography, next-generation IPv6 safeguards and de facto market standards/solutions used for credit card transactions appear to be the most significant technologies. Next-generation security must safeguard entire networks in addition to individual users, clients and servers.

Security is a moving target that will be solved mostly through independent efforts that eventually gain acceptance on the Internet. Security validation for virtual environments will hinge on the effectiveness of these methods when used together with challenging high-bandwidth fast-response applications.

- Economic problems on both personal and corporate scales often result from misunderstanding or misinformation. This happens everywhere, including the computing and academic communities. Marketing hyperbole takes advantage of the dizzying effects of ceaseless rapid change and the individual's need to stay informed. Rapidly growing bodies of knowledge and information are now part of the intellectual marketplace. Good and bad untested ideas go into the market early hoping to attract resources and influence the development of related work. However, the bottom line is results. Implementation (can it be built?), test (what does it really do?) and evaluation (does it do what we want?) are truer measures of success. Implementation, test and evaluation are the forcing functions that will produce working large-scale internetworked graphics environments. Visible results cut through the inevitable fog of hype and distorted information that surrounds rapid technical progress.

6. Applications. Working applications will drive progress, rather than theories or hype. In this section we present feasible applications that are exciting possibilities or existing works in progress. Many new projects are possible and likely in the near future if virtual environment requirements are adequately supported by the global information infrastructure. Several conjectured scenarios follow.

Sports: live 3D stadium with instrumented players. Imagine that all of the ballplayers in a sports stadium wear a small device that senses location (through the Global Positioning System and local spread-spectrum radio signal sensing) and transmits behavior packets over a wireless network. Similar sensors are embedded in gloves, balls, bats and even shoes. A computer server in the stadium feeds player telemetry into a physically based articulated human model, that in turn extrapolates individual body and limb motions. The server also maintains a scene database for the stadium complete with texture-mapped images of the edifice, current weather and representative pictures of fans in the stands. Meanwhile, Internet users have 3D browsers that can navigate and view the action from any perspective in the stadium. Users can also tune to multicast channels providing updated player positions and postures (via light-weight behavior interactions) along with live audio and video. Statistics, background information and multimedia home pages are available for each player. Some remote spectators might even choose to publicly share their own play-by-play commentary. Thus any number of remote fans might supplement traditional television coverage with a live interactive computer-generated view. Perhaps the most surprising aspect of this scenario is that all component software and hardware technologies exist today. A proposed personal tracking system (PTS) suitable for multiple players in an stadium environment is presented in Figure 6 [Bible 95].

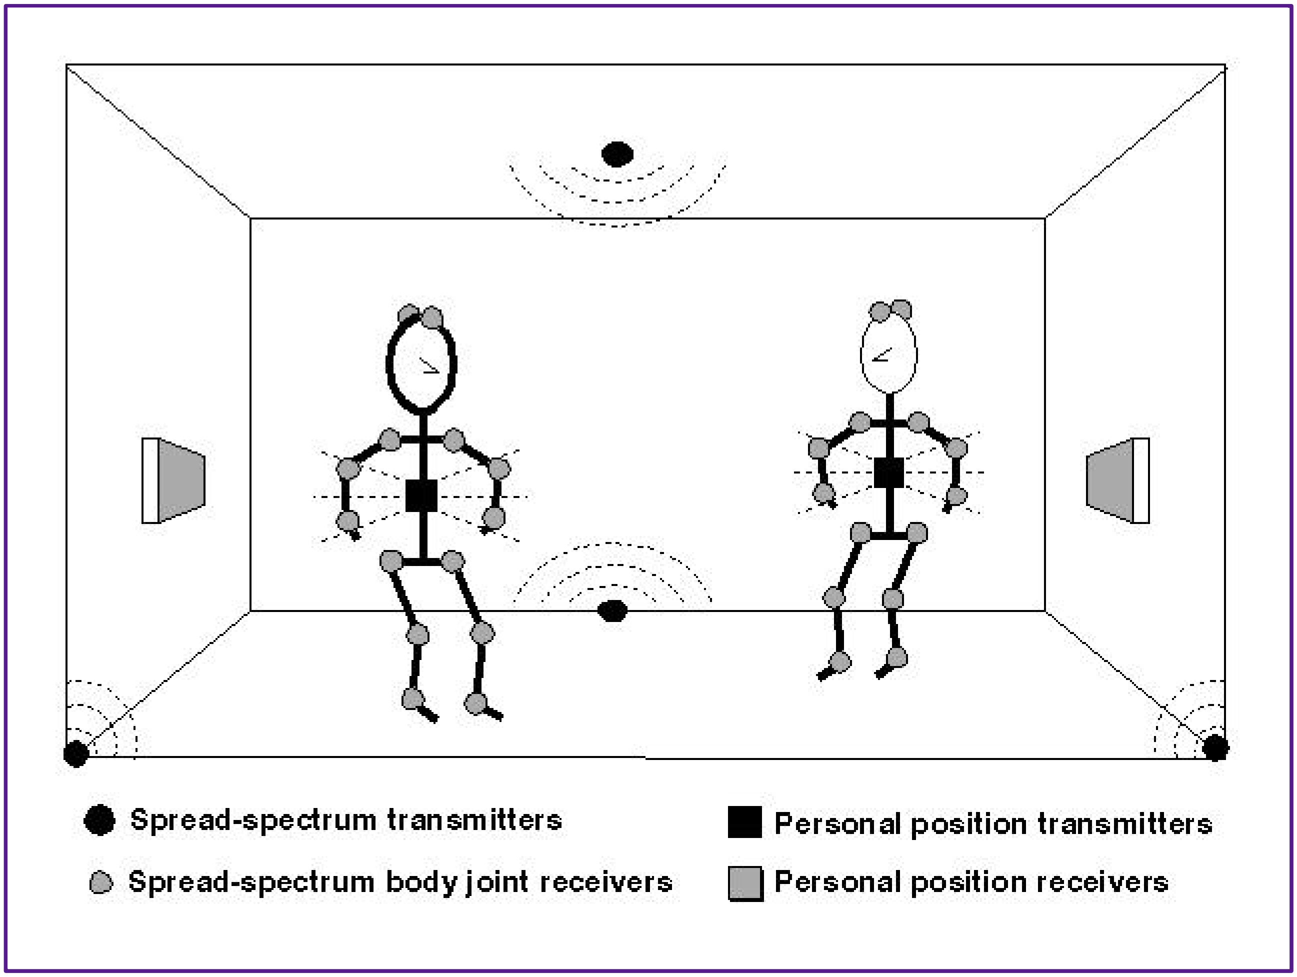

Figure 6. Example personal tracking system (PTS) using spread-spectrum radio ranging in a virtual environment [Bible 95]. Fourteen independent joint trackers passively determine body posture using a scheme similar to the Global Positioning System (GPS), then combined joint signals are multiplexed and transmitted back to the surrounding environment.

Education. Many people are trying to capitalize on the potential of global content by combining internetworking with education, long-distance teletraining and life-long learning. Applying computer graphics to education is also a steadily growing area of activity [Owen 95]. Our approach to learning and research includes grounding all theoretical concepts in real-world applications and events. Networked education efforts in the Monterey Bay region aim at enabling teachers and students of all ages to use all aspects of the Internet [Brutzman 95a, b] [Bigelow 95]. Since students study a great deal of content, regularly produce context and are excellent evaluators of effectiveness, it's likely that the best new examples combining innovative content with networked graphics will emerge from the educational community. Supporting educational network construction and teacher training is an excellent way to build a collaborative user base for new applications.

Military: 100,000 player problem. "Exploiting Reality with Multicast Groups" describes ground-breaking research on increasing the number of active entities within a virtual environment by several orders of magnitude [Macedonia 95a, b] Multicast addressing and the DIS protocol are used to logically partition network traffic according to spatial, temporal and functional entity classes as shown in Figure 7 . "Exploiting Reality" further explains virtual environment network concepts and includes experimental results. This work has fundamentally changed the distributed simulation community, showing that very large numbers of live and simulated networked players in real-world exercises are feasible. The military origins of this problem will become irrelevant as our abilities to effectively interact socially, scientifically and educationally scale to the hundreds of thousands. Such growth from military to real-world problems is analogous to the original design of the ARPANET and military predecessors to today's Internet. Many different types of simulations can be connected via distributed virtual environments. An excellent reference that describes and organizes a thorough variety of analytic simulation approaches is [Fishwick 95].

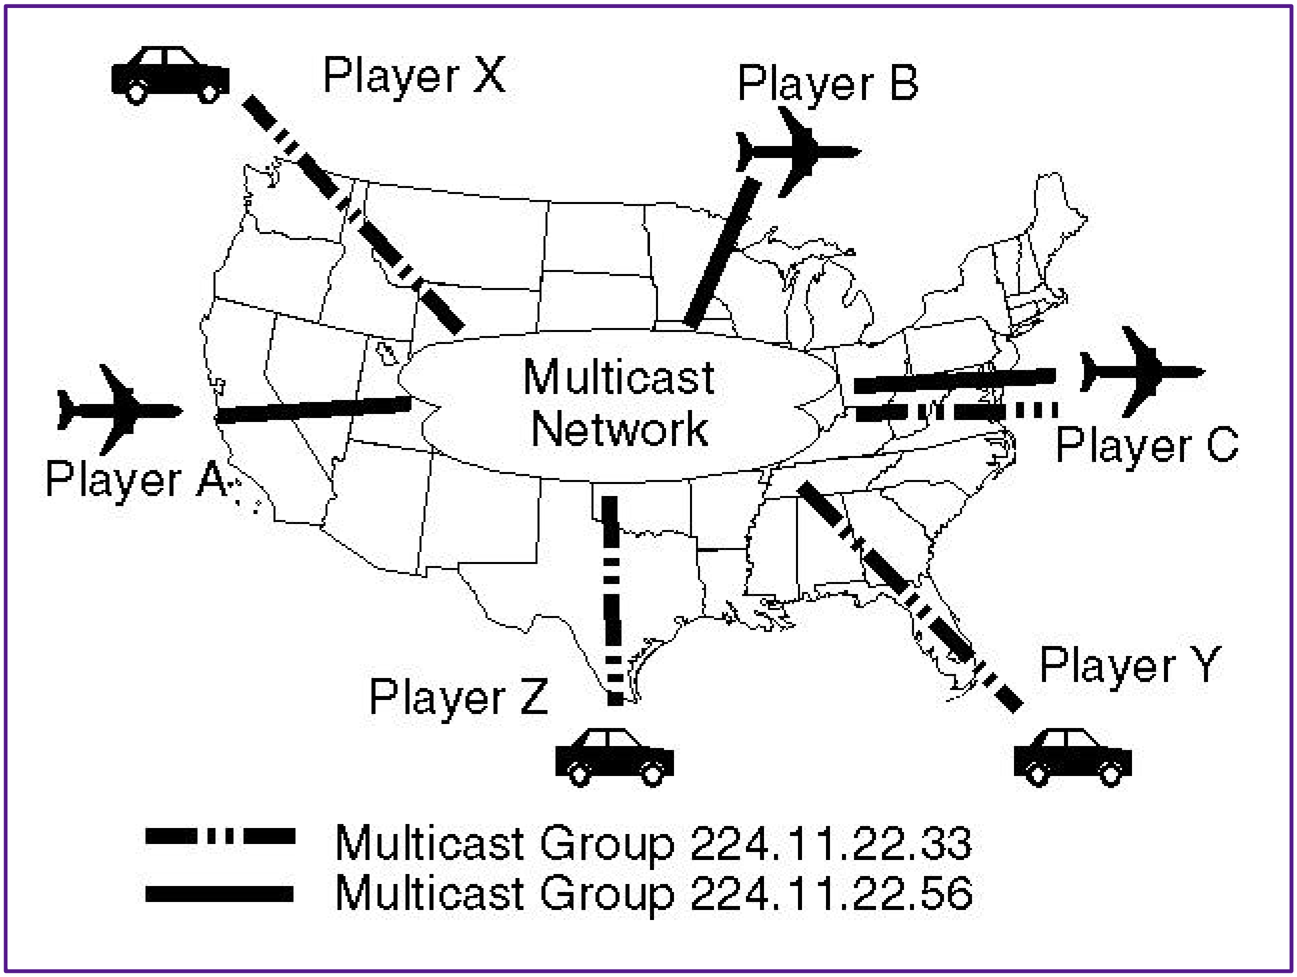

Figure 7. Multicast channels can partition network communications into spatial, temporal and functional classes to minimize traffic and exploit reality [Macedonia 95a, b] .

Science: virtual worlds as experimental laboratories for robots and people. In separate work, we've shown how an underwater virtual world can comprehensively model all salient functional characteristics of the real world for an Autonomous Underwater Vehicle (AUV) in real time [Brutzman 94]. This virtual world is designed from the perspective of the robot, enabling realistic AUV evaluation and testing in the laboratory. In this context, our "Turing Test" for the virtual environment is whether the robot behaves identically in the underwater virtual world and in the real world. Real-time 3D computer graphics are our window into that virtual world (Figure 8). Visualization of robot interactions within a virtual world permits sophisticated analyses of robot performance that are otherwise unavailable. Sonar visualization permits researchers to accurately "look over the robot's shoulder" or even "see through the robot's eyes" to intuitively understand sensor-environment interactions. Theoretical derivation of six-degree-of-freedom hydrodynamics equations has provided a general physics-based model capable of replicating highly nonlinear (yet experimentally verifiable) response in real time. Distribution of underwater virtual-world components enables scalability and rapid response. Networking allows remote access, demonstrated via MBone audio and video collaboration with researchers at distant locations. Integrating the World Wide Web allows rapid access to resources distributed across the Internet. Distributed resources include a text-to-speech sound server that translates robot reports into sound files [Belinfante 94]. Ongoing challenges include scaling up the types of interactions, datasets, models and live streams that can be coordinated within the virtual world.

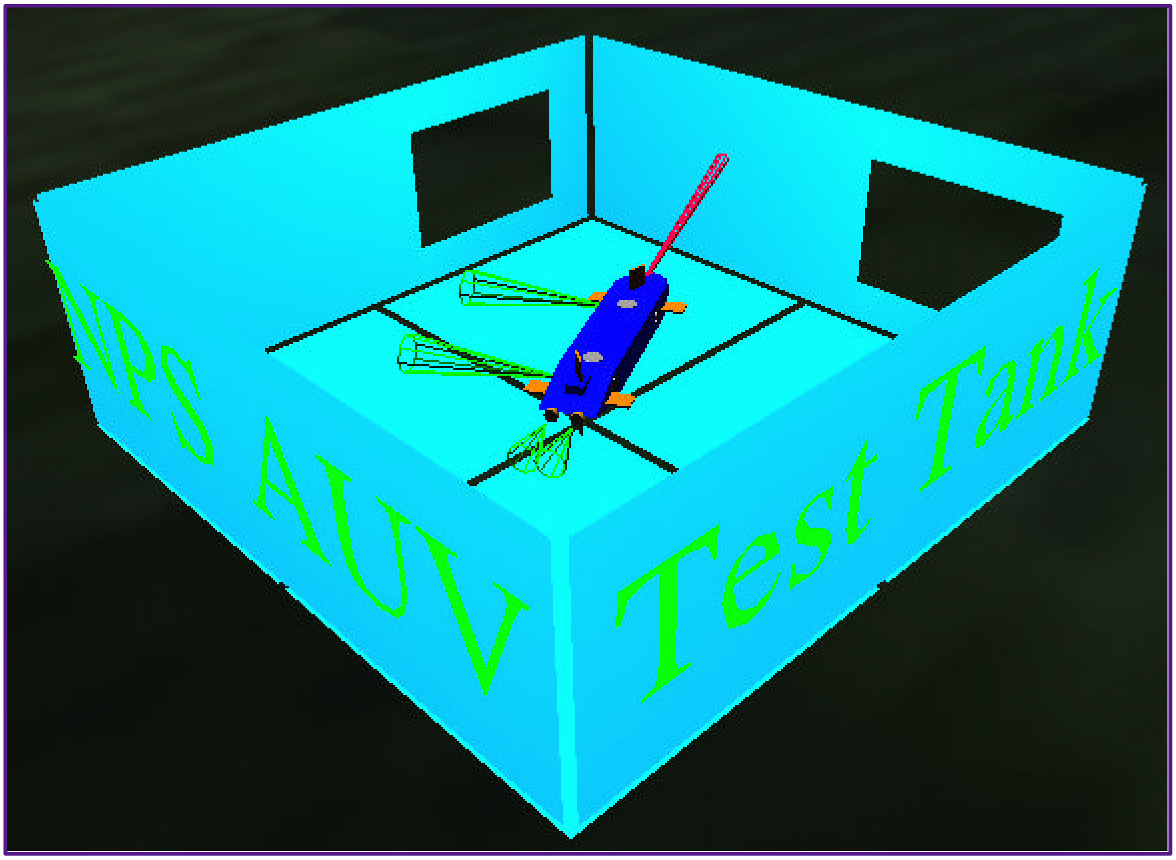

Figure 8. An autonomous underwater vehicle (AUV) operating in a virtual test tank. Virtual worlds can incorporate datasets, databases, models, robots and people [Brutzman 94].

Interaction: Multiple CAVEs using ATM and VRML. A CAVE is a walk-in synthetic environment that replaces the four walls of a room with rear-projection screens, all driven by real-time 3D computer graphics [Cruz-Neira 93]. These devices can accommodate 10-15 people comfortably and render high-resolution 3D stereo graphics at 15 Hz and higher update rates. The principal costs of a CAVE are in high-performance graphics hardware. We wish to demonstrate affordable linked CAVEs for remote group interaction. The basic idea is to send graphics streams from a master CAVE through a high-speed low-latency ATM link to a less expensive slave CAVE that contains only rear-projection screens.

Several serious ATM bottlenecks remain that must be resolved experimentally. Automatic generation of VRML scene graphs and simultaneous replication of state information over standard multicast links will permit both CAVEs and networked computers to interactively view results generated in real time by a supercomputer. Our initial application domain is a gridded virtual environment model of the oceanographic and biological characteristics of Chesapeake Bay, as shown in Figure 9 [Wheless 95, 96] [I-WAY 95]. To better incorporate networked sensors and agents into this and other virtual worlds, we are also investigating extensions to IP using underwater acoustics [Brutzman 95d]. As a final component, we are helping establish ambitious regional education and research networks which connect scientists, students from kindergartens through universities, libraries and the general public. Vertically integrated Web and MBone applications and the common theme of live networked environmental science are expected to provide many possible virtual-world connections [Brutzman 95a, b, c].

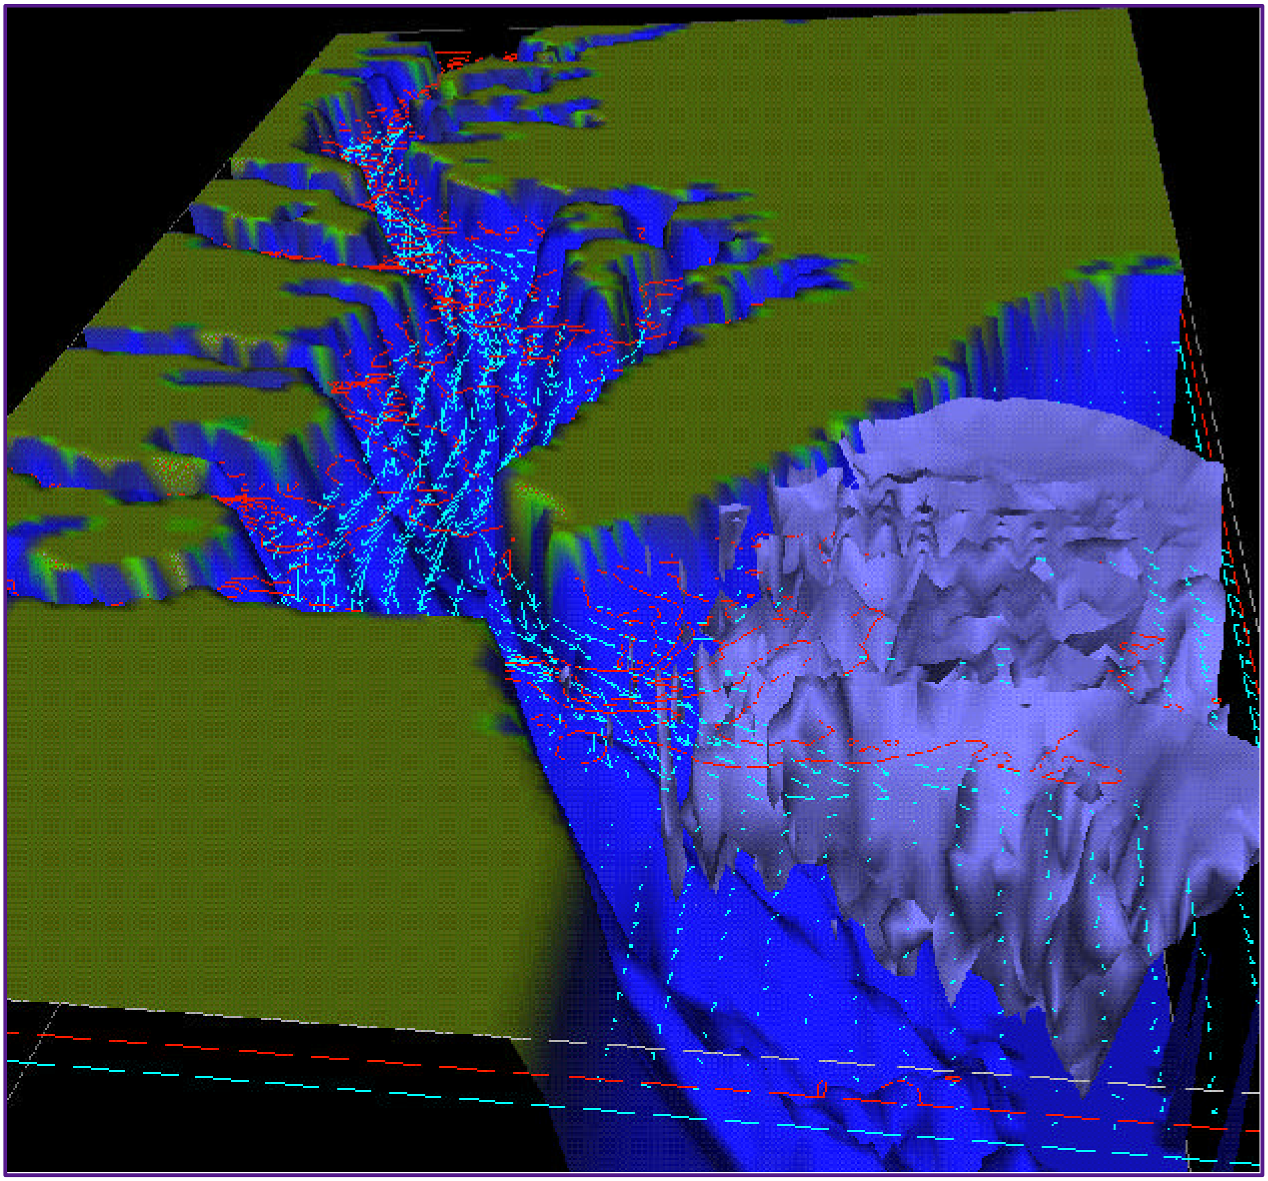

Figure 9. The Chesapeake Bay Virtual Environment (CBVE) combines multiple models, datasets and interaction modes with Sstandard IP and high-bandwidth ATM internetworking [Wheless 95, 96].

Mobile global public domain television. Access to the MBone means that anyone can become a global television provider. During SIGGRAPH 95 we experimented with mobile audio/video delivery to demonstrate location-independent multicasting. We loaded a cart with a camera-equipped SGI Indy workstation, a 2 Mbps wireless bridge, wireless microphones and video gear. Dubbed "MBone Unplugged," we were able to roam the convention floor while continuously multicasting exhibits and events world-wide (Figure 10) [Brutzman, Emswiler 95] [Clinger 96]. Working under conference auspices provides access to significant content with appropriate release permissions. We now believe that this rig is the workstation of the future: totally mobile, complete graphics and video capabilities, no apparent limits. Related work showed that live world-wide multicast of an academic course was possible for an entire quarter [Emswiler 95]. Planned future work includes simultaneous world-wide conference multicasting over ATM and MBone links while directly recording the audio/video at multiple bandwidths to a digital archive. We expect to show that previously transitory content produced by conference events can provide significant long-term value through easy, inexpensive live multicast and storage of high-bandwidth audiovisual streams.

Figure 10. "MBone Unplugged" cart used at SIGGRAPH 95 for wireless audio, video and 3D graphics with world-wide multicast scope [Brutzman, Emswiler 95].

---

7. Personal impacts. Tremendous enthusiasm and deep insights result when people use good interactive 3D graphics applications. Combining graphics and networks looks like a sure win, but the personal impacts of new technology are notoriously hard to predict. Here are a few ideas gained during our journey about how people will use large-scale internetworked virtual environments.

Large-scale means everyone and everything. This is a people process just as much as a technical process. Here are some people-related lessons learned from the ongoing growth of the Web. When something is easy, people do it. When something is challenging, people are willing to pay someone else or figure out another way to do it. When something is difficult, people don't do it. Examples of relatively easy things include building personal web pages and putting information online. Challenging things include professional web server providers and current web-server-in-a-box products. Excessively difficult things include early false starts by some commercial online services and overcomplex standardization efforts that never reach critical mass.

The personal impact of scaling up virtual environments to arbitrarily large sizes means that connecting and contributing to them can't be harder than setting up your own home page. If the process is not well defined, few people will risk wasting their time. As long as hooking up and authoring virtual worlds remain difficult, the technical solutions aren't yet good enough.

Information overload is the new norm. There's already too much information available. The unexpected will continue to happen, usually right when people think that they're caught up with what's current. The possibility of 500 television channels is ridiculously simple; there are already far more content sources on the Internet. With interactive opportunities increasing exponentially, personal impacts will likely have to be experienced to be understood. What does it mean when anyone can create, join, ignore or compete with any channel of information? No one really knows but we'll undoubtedly find out. Some individuals look to science fiction for inspiration and models of behavior. Some virtual reality architects on the VRML mailing list appear to go beyond inspiration and use Snow Crash [Stephenson 92] as a design document. Life imitates art, science imitates science fiction.

8. Projections. Merely by reading the New York Times, any individual can get more information about the world than was available to any world leader throughout most of human history. Multiply that single stream by the millions of other information sources on the Internet, and it's clear we do not lack content. Mountains of content are accessible. What we need more is context, i.e. a way to interactively locate, retrieve and display related pieces of information in a timely manner. Context establishes the reliability and pedigree of information streams. It also permits selectively culling large volumes of unstructured information that can otherwise overwhelm any single person or process.

Within two lifetimes we have invented several new technologies for recording and exchanging information. Handwriting gave way to typing, then typing to word processing, and soon afterward came desktop publishing. Now people can use 3D real-time interactive graphics simulations and dynamic "documents" with multimedia hooks to record and communicate information. Such "documents" can be directly distributed on demand to anyone connected to the Internet. In virtual environments a further paradigm shift becomes possible. The long-term potential of virtual environments is to serve as an archive and interaction medium, combining massive and dissimilar datasets and data streams of every conceivable type. Virtual environments will enable comprehensive and consistent interaction within those massive datasets, data streams and models that recreate reality. Virtual environments can provide meaningful context to the mountains of content that currently exist in isolation without roads, links or order.

What about scaling up? Fortunately there already exists a model for these growing mountains of information content: the real world. Virtual worlds address the context issue by providing information links similar to those in the real world. When our virtual constructs cumulatively approach realistic levels of depth and sophistication, our understanding of the real world will deepen correspondingly. In support of this goal, we have shown how the structure and scope of virtual environment relationships can be dynamically extended using feasible network communications methods. This efficient distribution of information will let any remote user or entity in a virtual environment participate and interact in increasingly meaningful ways.

Open access to any type of live or archived information resource is becoming available for everyday use by individuals, programs, collaborative groups and even robots. Virtual environments are a natural way to provide order and context to these massive amounts of information. World-wide collaboration works, for both people and machines. Finally, the network is more than a computer, and even more than your computer. The Internet becomes our computer as we learn how to share resources, collaborate and interact on a global scale.

9. Conclusions. Interactive 3D computer graphics enables people to see things that are otherwise imperceptible, impractical or impossible. Computer networking can connect everyone to everything, any person to any machine to any information resource. Networks are not "something else" relative to graphics, rather they are integral as we scale up in every direction. As these interrelated technologies combine, understanding why limits exist will help us break through bottlenecks. Each bottleneck broken reveals another as we consider new-found capabilities and their corresponding limitations.

Large-scale virtual environments are the new grand challenge in computing, a challenge that can include every person and every subject. The best motivation for scaling up can be found in the respective strengths of 3D graphics and networking: perception and bandwidth. Interactive graphics let us perceive the world in new and familiar ways. Networks provide information bandwidth, broad connectivity and varied content at rapid rates. Our new challenges are fascinating. Combining graphics and networks now lets us work on perceptual bandwidth, maximizing the depth and breadth and speed we use to interpret the world around us. This achievement may be our biggest breakthrough.

References.

Belinfante, Axel, "Say..." text to speech sound server, University of Twente, Netherlands, 1994. Available at *http://wwwtios.cs.utwente.nl/say/?*

Bell, Gavin, Parisi, Anthony and Pesce, Mark, "The Virtual Reality Modeling Language (VRML) Version 1.0 Specification," 26 May 1995. Available via the VRML Repository at *http://www.sdsc.edu/vrml*

Berners-Lee, Tim, Cailliau, Luotonen, Ari, Nielsen, Henrik Frystyk, and Secret, Arthur, "The World Wide Web," *Communications of the ACM*, vol. 37 no. 8, August 1994, pp. 76-82.

Bible, Steven R., Zyda, Michael and Brutzman, Don, "Using Spread-Spectrum Ranging Techniques for Position Tracking in a Virtual Environment," *IEEE Networked Realities Workshop*, Boston Massachusetts, October 26-28 1995. Available at *http://www-npsnet.cs.nps.navy.mil/npsnet/publications/NR95-Paper-Bible.ps.Z*

Bigelow, Randall J., *Internetworking: Planning and Implementing a Wide-Area Network (WAN) for K-12 Schools*, Master's Thesis, Naval Postgraduate School, Monterey California, June 95. Available at *http://www.stl.nps.navy.mil/~rjbigelo/thesis.html*

Bradner, Scott O. and Mankin, Allison, editors, *IPng: Internet Protocol Next Generation*, Addison-Wesley, Reading Massachusetts, 1996. Additional information available at *http://playground.sun.com/pub/ipng/html/ipng-main.html*

Brutzman, Donald P., *A Virtual World for an Autonomous Underwater Vehicle*, Ph.D. Dissertation, Naval Postgraduate School, Monterey California, December 1994. Available at *http://web.nps.navy.mil/~brutzman/dissertation/*

Brutzman, Don, "Networked Ocean Science Research and Education, Monterey Bay California," *Proceedings of International Networking (INET) 95 Conference*, Internet Society, Honolulu Hawaii, June 27-30 1995. Available at *http://www.isoc.org/HMP/PAPER/039/abst.html*

Brutzman, Don, "Remote Collaboration with Monterey Bay Educators," *Visual Proceedings, Association for Computing Machinery (ACM) Special Interest Group on Computer Graphics (SIGGRAPH) 95*, Los Angeles California, August 7-11 1995, p. 145.

Brutzman, Don and Emswiler, Tracey, "MBone Unplugged," SIGGRAPH 95 exhibit report, August 1995. Available at *http://web.nps.navy.mil/~brutzman/unplugged.html*

Brutzman, Don and Reimers, Stephen, "Internet Protocol over Seawater: Towards Interoperable Underwater Networks," *Unmanned Untethered Submersibles Technology 95*, Northeastern University, Nahant Massachusetts, September 25-27 1995, pp. 444-457. Available at *ftp://taurus.cs.nps.navy.mil/pub/auv/ipoversw.ps*

Carey, Rikk, Marrin, Chris and Bell, Gavin, editors, "The Virtual Reality Modeling Language (VRML) Version 2.0 Specification," International Standards Organization/International Electrotechnical Commission (ISO/IEC) draft standard 14772, August 4 1996. Available via the VRML Repository at *http://www.sdsc.edu/vrml*

Clinger, Marke, "GraphicsNet '95: Integrated Voice, Video, Graphics and Data Network Using Asynchronous Transfer Mode (ATM)," *COMPUTER GRAPHICS*, ACM SIGGRAPH, vol. 30. no. 1, February 1996, pp. 10-18. Additional information at http://www.siggraph.org/conferences/siggraph95/GraphicsNet

Cogger, Dick, CU-SeeMe, desktop videoconferencing software, Cornell University, Ithaca New York, 1995. Available at *ftp://gated.cornell.edu/pub/video/html/Welcome.html*

Comer, Douglas E., *Internetworking with TCP/IP Volume I: Principles, Protocols and Architecture*, second edition, Prentice Hall, Englewood Cliffs New Jersey, 1991.

Cooke, J.C., Zyda, M.J., Pratt, D.R. and McGhee, R.B., "NPSNET: Flight Simulation Dynamic Modeling using Quaternions," *PRESENCE: Teleoperations and Virtual Environments*, vol. 1 no. 4, Fall 1992, pp. 404-420. Available at
http://www-npsnet.cs.nps.navy.mil/npsnet/publications/NPSNET.Flight.Simulation.Dynamic.Modeling.Using.Quatern

Cruz-Neira, Carolina, Leigh, Jason, Papka, Michael, Barnes, Craig, Cohen, Steven M., Das, Sumit, Engelmann, Roger, Hudson, Randy, Roy, Trina, Siegel, Lewis, Vasilakis, Christina, DeFanti, Thomas A. and Sandin, Daniel J., "Scientists in Wonderland: A Report on Visualization Applications in the CAVE Virtual Reality Environment," *IEEE 1993 Symposium on Research Frontiers in Virtual Reality*, San Jose California, October 25-26 1993, pp. 59-66 and CP-3.

Curtis, Pavel and Nichols, David A., "MUDs Grow Up: Social Virtual Reality in the Real World," *Proceedings of the IEEE Computer Conference*, IEEE Computer Society Press, Los Alamitos California, 1994, pp. 193-200. Available at *ftp://ftp.parc.xerox.com/pub/MOO/papers/MUDsGrowUp.ps*

Deering, Steve, "Host Extensions for IP Multicasting," Request for Comments (RFC) 1112, Internet Engineering Task Force (IETF), August 1989. Available at *ftp://ds.internic.net/rfc/rfc1112.txt*

DIS Steering Committee, *The DIS Vision: A Map to the Future of Distributed Simulation*, version 1, May 1994. Available at *ftp://sc.ist.ucf.edu/public/STDS/vision.ps*

Durlach, Nathaniel I. and Mavor, Anne S., editors, *Virtual Reality: Scientific and Technological Challenges*, National Research Council, National Academy Press, Washington DC, 1995.

Emswiler, Tracey, *Internetworking: Using the Multicast Backbone (MBone) for Distance Learning*, Master's Thesis, Naval Postgraduate School, Monterey California, September 95. Summary video available at *http://www.stl.nps.navy.mil/~iirg/emswiler/emswiler.qt.Z*

Fishwick, Paul A., *Simulation Model Design and Execution: Building Digital Worlds*, Prentice Hall, Englewood Cliffs New Jersey, 1995. Additional information and *SimPack* simulation toolkit software available at *http://www.cis.ufl.edu/~fishwick/*

Garrett, John and Waters, Donald, cochairs, "Preserving Digital Information," Task Force on Archiving of Digital Information, Commission on Preservation & Access and the Research Libraries Group, draft report version 1.0, Yale University, New Haven Connecticut, August 24 1995.

Gelernter, David, *Mirror Worlds - or the Day Software Puts the Universe in a Shoebox... How It Will Happen and What It Will Mean*, Oxford University Press, New York, 1992.

Hughes, Kevin, "Entering the World Wide Web (WWW): A Guide to Cyberspace," Enterprise Integration Technology Inc., May 1994. Available at *http://www.eit.com/web/www.guide/*

IEEE Standard for Information Technology - Protocols for Distributed Interactive Simulation (DIS) Applications, version 2.0, Institute for Simulation and Training report IST-CR-93-15, University of Central Florida, Orlando Florida, May 28 1993.

International Wide-Area Year (I-WAY) project, North American ATM network in support of *IEEE/ACM Supercomputing 95*, San Diego California, December 3-7 1995. Information available at *http://www.iway.org*

Internet Network Information Center (Internic), Request For Comments (RFC) archive, *ftp://ds.internic.net*, 1995. Hypertext interface and search facilities available at *http://www.internic.net*

Library of Congress, home page, November 1995. Available at *http://www.loc.gov/*

Macedonia, Michael R. and Brutzman, Donald P., "MBone Provides Audio and Video Across the Internet," *IEEE COMPUTER*, April 1994, pp. 30-36. Available at *ftp://taurus.cs.nps.navy.mil/pub/i3la/mbone.html*

Macedonia, Michael R., *A Network Software Architecture for Large Scale Virtual Environments*, Ph.D. Dissertation, Naval Postgraduate School, Monterey California, June 1995. Available at *http://www.cs.nps.navy.mil/research/npsnet/publications/Michael.Macedonia.thesis.ps.Z*

Macedonia, Michael R., Zyda, Michael J., Pratt, David R., Brutzman, Donald P. and Barham, Paul T., "Exploiting Reality with Multicast Groups: A Network Architecture for Large-Scale Virtual Environments," *IEEE Computer Graphics and Applications*, vol. 15 no. 5, September 1995, pp. 38-45. Available at *http://www.cs.nps.navy.mil/research/npsnet/publications/IEEECGA.ps.Z*

Malamud, Carl, *Stacks: Interoperability in Today's Computer Networks*, Prentice Hall, Englewood Cliffs New Jersey, 1992.

Netscape Corporation, "An Exploration of Dynamic Documents," online documentation page, November 1995. Available at *http://www.netscape.com/assist/net_sites/pushpull.html*

NRENAISSANCE Committee, *Realizing the Information Future: The Internet and Beyond*, Computer Science and

Telecommunications Board, National Research Council, National Academy Press, Washington DC, 1994. Available at *http://www.nap.edu/nap/online/rtif/*

Ousterhout, John K., *Tcl and the Tk Toolkit*, Addison-Wesley, Reading Massachusetts, 1994.

Owen, Scott, guest editor, focus issue on computer graphics education, *COMPUTER GRAPHICS, ACM SIGGRAPH*, vol. 29 no. 3, August 1995.

Pesce, Mark and Behlendorf, Brian, moderators, "Virtual Reality Modeling Language (VRML)," working group mail list, 1994-1996. Archived at *http://vrml.wired.com/* and *http://www.sdsc.edu/vrml*

Pesce, Mark, *VRML - Browsing and Building Cyberspace*, New Riders Publishing, Indianapolis Indiana, 1995.

Schulzrinne, Henning, Casner, Stephen, Frederick, Ron and Jacobson, Van, "RTP: A Transport Protocol for Real-Time Applications," Request for Comments (RFC) 1889, Internet Engineering Task Force (IETF), January 1996. Available at *ftp://ds.internic.net/rfc/rfc1889.txt*

Stallings, William, *Data and Computer Communications*, fourth edition, Macmillan Publishing, New York, 1994.

Stephenson, Neal, *Snow Crash*, Bantam Books, New York, 1992.

Stevens, Richard W., *UNIX Network Programming*, PTR Prentice Hall, Englewood Cliffs New Jersey, 1990.

Stone, Steven, *NPSNET: A Rapidly Reconfigurable Application-Layer Virtual Environment Network Protocol*, Master's Thesis, Naval Postgraduate School, Monterey California, September 1996.

Sun Microsystems Corporation, Java language home page, November 1995. Available at *http://java.sun.com/*

Wernicke, Josie, T*he Inventor Mentor: Programming Object-Oriented 3D Graphics with OpenInventorTM*, Release 2, Addison-Wesley Publishing, Reading Massachusetts, 1994. Summary available at *http://www.sgi.com/Technology/Inventor/VRML/TIMSummary.html*

Wheless, Glen, LaScara, Cathy, Valle-Levison, Arnaldo, Brutzman, Don and Sherman, Bill, "Chesapeake Bay Virtual Ecosystem Model (CBVEM): Interacting with a Couple Bio-Physical Simulation," *IEEE/ACM Supercomputing 95*, San Diego California, December 3-7 1995. Information available at *http://www.ccpo.odu.edu/~wheless*

Wheless, Glen H., LaScara, Cathy M., Valle-Levison, Arnaldo, Brutzman, Donald P., Sherman, William, Hibbard, William L. and Paul, Brian E., "Chesapeake Bay Virtual Ecosystem Model (CBVEM): initial results from the prototypical system," *International Journal of Supercomputer Applications*, to appear. Available at ftp://ftp.ccpo.odu.edu/pub/wheless/wheless_ijsa.tar

White, James E., Telescript Technology: Mobile Agents, white paper, General Magic Inc., Sunnyvale California, 1995. To appear in Bradshaw, Jeffrey, ed., *Software Agents*, AAAI Press/MIT Press, Menlo Park California, 1996.

Zyda, Michael J., Pratt, David R., Falby, John S., Barham, Paul T., Lombardo, Chuck and Kelleher, Kristen M., "The Software Required for the Computer Generation of Virtual Environments," *PRESENCE: Teleoperators and Virtual Environments*, vol. 2 no. 2, MIT Press, Cambridge Massachusetts, Spring 1993, pp. 130-140. Available at http://www-npsnet.cs.nps.navy.mil/npsnet/publications/NPSNET.Software.Required.for.the.Computer.Generation.of.

Acknowledgements. I gratefully acknowledge contributions and insights gained from discussions with Mike Macedonia, Mike Zyda, Glen Wheless, Clark Dodsworth and members of the Virtual Reality Modeling Language (VRML) mail list. Funding support for this research was provided by the Naval Postgraduate School and the Office of Naval Research.

Author Information. Don Brutzman is a computer scientist working in the Interdisciplinary Academic Group at the Naval Postgraduate School. His research interests include underwater robotics, real-time 3D computer graphics, artificial intelligence, high-performance networking and large-scale virtual environments. He is a member of the Institute of Electrical and Electronic Engineers (IEEE), the Association for Computing Machinery (ACM) Special Interest Group on Graphics (SIGGRAPH), the American Association for Artificial Intelligence (AAAI), the Marine Technology Society (MTS) and the Internet Society (ISOC). Collaborative project information is available at

http://web.nps.navy.mil/~brutzman

The Uniform Resource Locator (URL) for this paper is
http://web.nps.navy.mil/~brutzman/vrml/breakthroughs.html
Your comments are welcome.

Maintainer: Don Brutzman (brutzman@nps.navy.mil)
Completed: 15 August 1997 Web links updated: 2 April 2001